

We claim:

1. A computer implemented method of processing a transaction to determine the risk of transaction, the method comprising:

storing a plurality of merchant clusters, the merchant clusters determined from statistical

5 co-occurrences of the merchant names in a plurality of transactions;

receiving data from a transaction between a consumer and merchant;

determining one of the plurality of merchant clusters associated with the merchant of the

transaction based on the merchant's name; and

applying the merchant cluster in conjunction with data derived from the transaction to a

10 predictive model to determine a level of risk of the transaction.

2. The method of claim 1, further comprising:

estimating a likelihood that the transaction is fraudulent.

3. The method of claim 1, responsive to the level of risk of the transaction, determining whether to approve the transaction, decline the transaction, or obtain additional information
15 regarding the transaction or the cardholder.

4. The method of claim 1, wherein determining one of the plurality of merchant clusters associated with the merchant further comprises:

storing a lookup table associating each merchant cluster with at least one merchant name,

wherein the merchant names are each unique;

20 applying the merchant's name to the lookup table to determine the associated merchant cluster.

5. The method of claim 4, wherein the unique merchant names are derived from a plurality of raw merchant names in transaction data by stemming and equivalencing the raw merchant names.

6. The method of claim 1, further comprising:

5 storing for each merchant cluster a risk factor indicative of the likelihood that transactions at merchants within the merchant cluster are fraudulent; and applying the risk factor of the merchant cluster to the predictive model.

7. The method of claim 6, wherein the risk factor is an estimate of the percentage of transactions in the merchant cluster that are fraudulent.

10 8. The method of claim 1, further comprising:

storing a plurality of consumer clusters;

storing for each combination of a consumer cluster and a merchant cluster a risk factor

indicative of the likelihood that transactions by consumers in the consumer cluster at merchants within the merchant cluster are fraudulent;

15 determining a current cardholder cluster associated with the cardholder; and

applying the risk factor of the combination of the current cardholder cluster and the merchant cluster to the predictive model.

9. The method of claim 8, wherein the risk factor is an estimate of the percentage of transactions in the merchant cluster by consumers in the cardholder cluster that are fraudulent.

20 10. The method of claim 1, further comprising:

storing for each merchant cluster at least one summarized transaction statistic,

descriptive of transactions occurring at merchants in the merchant cluster; and

applying the at least one summarized transaction statistic of the merchant cluster to the predictive model.

11. The method of claim 10, wherein the at least one summarized transaction statistic are selected from a group consisting of:

5 average transaction amount; and
 average transaction volume.

12. The method of claim 1, further comprising:

storing for each of a plurality of consumer clusters, at least one summarized transaction
statistic, descriptive of transactions by consumers in the consumer cluster; and
10 applying the at least one summarized transaction statistic of the consumer cluster to the
 predictive model.

13. The method of claim 10, wherein the at least one summarized transaction statistic are selected from a group consisting of:

15 average transaction amount; and
 average transaction volume.

14. A computer implemented method of processing a transaction to determine the risk of transaction, the method comprising:

storing a plurality of merchant clusters, the merchant clusters determined from statistical
co-occurrences of the merchant names in a plurality of transactions;

20 receiving data of a transaction between a consumer and merchant;

determining one of the plurality of merchant clusters associated with the merchant of the
transaction based on the merchant name;

determining an affinity measure of an affinity of cardholder to the merchant cluster; and

applying the affinity measure in conjunction with data derived from the transaction to a predictive model to determine the level of risk of the transaction.

15. The method of claim 6, wherein determining the affinity measure of an affinity of the cardholder to the merchant cluster further comprises:

5 determining an affinity vector of the affinity of the cardholder to each of a plurality of merchant clusters, including the merchant cluster of the merchant of the transaction.

16. The method of claim 6, wherein determining the affinity measure of an affinity of the cardholder to the merchant cluster further comprises:

10 determining a cardholder cluster associated with the cardholder; and
determining an affinity measure of the affinity of the cardholder cluster to the merchant cluster.

17. The method of claim 6, wherein determining the affinity measure of an affinity of the cardholder to the merchant cluster further comprises:

15 determining a cardholder cluster associated with the cardholder; and
determining an affinity vector of the affinity of the cardholder cluster to each of a plurality of merchant clusters, including the merchant cluster of the merchant of the transaction.

18. A method of determining the level of risk in a transaction by consumer, the method comprising:

20 storing a plurality of merchant clusters, the merchant clusters determined from statistical co-occurrences of the merchant names in a plurality of transactions;
receiving data of a current transaction between a consumer and merchant;

determining a predicted merchant cluster in which the consumer is predicted to have a future transaction based on transactions of the consumer prior to the current transaction;

determining an actual merchant cluster associated with the merchant of the transaction based on the merchant name;

determining a difference measure between the predicted merchant cluster and the actual merchant cluster; and

applying the difference measure in conjunction with data derived from the transaction to a predictive model to determine the level of risk of the transaction.

19. A system for detecting risk in a transaction, comprising:

a database of unique merchant names, each merchant name associated with a merchant cluster;

a transaction processing component that receives a transaction between a consumer and a merchant, that derives transaction data from the transaction, and determine a unique merchant name for the merchant from the database; and

a statistical model that receives the data derived from the transaction and the unique merchant name, and outputs a score indicative of the level of risk in a transaction.

20. A method of determining a level of risk in a transaction, the method comprising:

receiving a transaction between a first entity and a second entity;

deriving high categorical information elements from at least one of transaction, the first entity or the second entity;

determining a low categorical information cluster closest to the high categorical information elements;

applying the low categorical information cluster and data derived from the transaction to a predictive model to determine the level of risk in the transaction.

21. The method of claim 20, wherein the high categorical information elements are text data.

22. The method of claim 20, wherein the second entity is a merchant of the transaction, and the high categorical information elements are merchant notes associated with the transaction.

5 23. The method of claim 20, further comprising:

selecting a plurality of high categorical information elements;

associating each high categorical information element with a context vector in a vector space, such that high categorical information elements that frequently proximally co-occur in the transactions have context vectors that are similarly oriented in the vector space;

10 clustering the context vectors of the high categorical information elements into a number of clusters substantially less than the number of high categorical information elements, each cluster being a low categorical information cluster;

wherein determining a low categorical information cluster closest to the high categorical information elements further comprises determining the low categorical information cluster closest in the vector space to a context vector derived from the context vectors of the high categorical information elements.